Hamiltonian analogs of combustion engines: a systematic exception to adiabatic decoupling

Lukas Gilz, Eike Thesing, and James R. Anglin
State Research Center OPTIMAS and Fachbereich Physik,
Technische Univerität Kaiserslautern, D-67663 Kaiserslautern, Germany
(Dated: January 19, 2017)

Workhorse theories throughout all of physics derive effective Hamiltonians to describe slow time evolution, even though low-frequency modes are actually coupled to high-frequency modes. Such effective Hamiltonians are accurate because of *adiabatic decoupling*: the high-frequency modes 'dress' the low-frequency modes, and renormalize their Hamiltonian, but they do not steadily inject energy into the low-frequency sector. Here, however, we identify a broad class of dynamical systems in which adiabatic decoupling fails to hold, and steady energy transfer across a large gap in natural frequency ('steady downconversion') instead becomes possible, through nonlinear resonances of a certain form. Instead of adiabatic decoupling, the special features of multiple time scale dynamics lead in these cases to efficiency constraints that somewhat resemble thermodynamics.

arXiv:1701.05006v1 [physics.class-ph] 18 Jan 2017

I. INTRODUCTION

510.0pt

A. Adiabatic decoupling

One of the most basic and broadly applicable principles in all of physics is *adiabatic decoupling*. It is already familiar from school physics: two weakly coupled degrees of freedom will only exchange substantial amounts of energy if their resonant frequencies match. In more sophisticated language, adiabatic decoupling means that the effects of fast degrees of freedom on slow degrees of freedom can be described by renormalizing the effective Hamiltonian which governs the slow time evolution. Applications of this principle range from the Born-Oppenheimer approximation in molecular physics [1, 2] to effective Lagrangians and renormalization in quantum field theory [3–5]. Whether the fast degrees of freedom in question are gluons in nuclear matter, atomic excitations, or vibrations in a rigid body, it is thanks to adiabatic decoupling that we can 'adiabatically eliminate' or 'integrate out' high-frequency modes, to leave an effective Hamiltonian which describes only slow variables, and yet describes them accurately, even over long times.

While the power of adiabatic elimination is often celebrated, its most basic implication is less frequently noted: since the slow evolution is Hamiltonian, it conserves energy. Fast degrees of freedom cannot steadily inject energy into the low-frequency world. Or can they?

B. Steady downconversion

Consider the case of combustion engines. These are such large and complex dynamical systems that they must in practice be described thermodynamically, rather than in microscopic terms. In principle, however, a real engine could still operate within a sealed container, at least for some time, as long as the container were large enough to include air, exhaust, and so on. The operation of a combustion engine is thus in principle a case of closed-system Hamiltonian time evolution. What a combustion engine nonetheless does, however, is to extract energy steadily from chemical degrees of freedom, whose dynamical time scales are on the order of femtoseconds (frequencies $\sim 10^{14}$ Hz), and then ultimately use this energy to turn wheels at a few thousand revolutions per minute ($\sim 10^2$ Hz). If we introduce the term *steady downconversion* to mean this kind of steady transfer of energy across a dynamical frequency gap, precisely as is normally forbidden by adiabatic decoupling, then steady downconversion across a frequency ratio of 10^{12} is in fact achievable, albeit in dynamical systems too large and complex for microscopic description.

C. Microscopic models

The question of exactly how it is that macroscopic engines manage to evade adiabatic decoupling, and achieve steady downconversion, leads quickly into the general question of how thermodynamics 'emerges' from mechanics in the limit of large system size. Much has been learned about that emergence, however, by starting from the microscopic side of the limit. The main goal of such work has hitherto been to learn how the ensembles of statistical mechanics can be justified by effective ergodization within Hamiltonian mechanics. As a fine example of this kind of study, we are inspired by the classic work of Chirikov, in showing how a not-too-small perturbation of an integrable system can break the invariant tori of Kolmogorov-Arnold-Moser theory [6] and lead to effectively ergodic behavior [7].

In particular, Chirikov presented instructive models of nonlinear oscillation, and showed how certain resonant perturbations can induce Arnold diffusion [7], even in

systems with few degrees of freedom. In the present paper we take a similar approach, to the point of studying some quite similar Hamiltonians. The significant differences that we will introduce, however, will reflect a quite orthogonal perspective: our concern is not information, but energy. Instead of aiming to understand the effective loss of information that is represented in statistical mechanics by ensembles, we will focus on steady downconversion as a mechanical analog to the conversion of energy-dense fuel into work, as in a combustion engine. Our basic question is therefore: In a phase space of only a few dimensions, what kind of Hamiltonian system can exhibit steady downconversion?

D. Hamiltonian 'daemons'

Because our central concern is thus secular energy transfer between slow and fast degrees of freedom, it is an important restriction that we will further insist on a *time-independent* Hamiltonian, so that the system conserves its total energy and no important energetic behavior is 'put in by hand'. This restriction to time-independent Hamiltonians means that although we are only seeking a sort of minimal toy model for a combustion engine, yet in an important way our quest is more ambitious than most theoretical models for microscopic engines: we do not allow the kind of macroscopic elements that are represented by external control parameters (which would make the Hamiltonian time-dependent) or by heat reservoirs (which would make the evolution non-Hamiltonian).

The kind of system we seek can thus be thought of as a sort of mechanical Maxwell's Demon — except that instead of being directed against the Second Law of Thermodynamics, the principle it contrives to violate is that of adiabatic decoupling. It is also an important difference that our small Hamiltonian systems constitute our entire systems of study, and are not, like Maxwell's Demon, small additions to macroscopically large systems. Insofar as we are considering microscopic analogs of engines, we are considering *entirely* microscopic analogs, not just microscopic working parts which extract energy from macroscopic heat baths. In our discussion at the end of this paper we will briefly compare this kind of isolated small system with systems coupled to baths.

To suggest the kinship of our systems with the Demon, while also acknowledging the important differences, we will refer more briefly to "small dynamical systems with time-independent Hamiltonians which exhibit steady energy transfer from fast to slow motion" as *Hamiltonian daemons*. The first question of our paper is then: Do daemons exist, mathematically? The answer can be seen from Fig. 1.

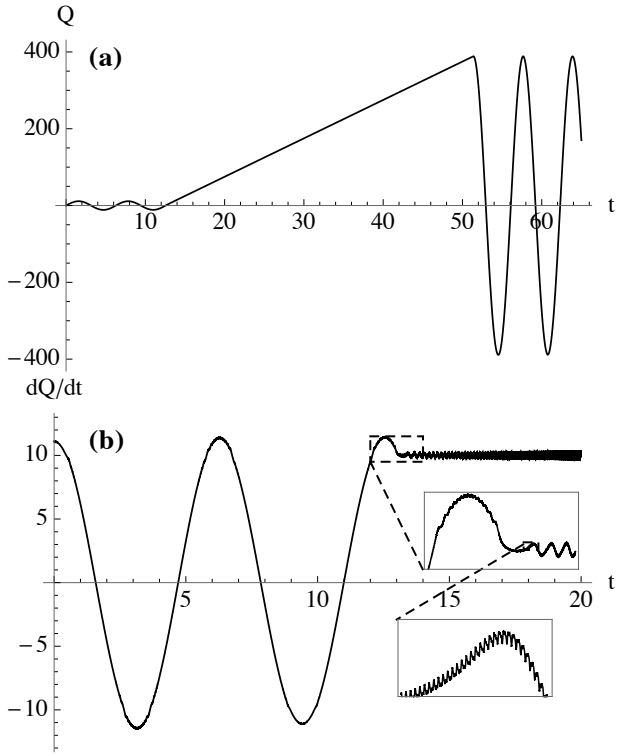

FIG. 1. Evolution of a Hamiltonian daemon, under Hamiltonian H_1 of Eqn. (1), with the parameter values stated in the text, and initial conditions $Q(0) = 0, \dot{Q}(0) = 11.1, q_+(0) = 40/\sqrt{15}, q_-(0) = p_+(0) = p_-(0) = 0$. **(a)** Position $Q(t)$ of a mass. **(b)** Velocity $\dot{Q}(t)$, with insets showing detail.

E. An example of a daemon

What Fig. 1 shows is the position Q (upper plot) and velocity $\dot{Q}$ (lower plot) of a certain point mass, which moves in time (horizontal axes) under the canonical equations of motion generated by the Hamiltonian (1), which we will present in Section II and then analyze for the bulk of this paper. Two other oscillators are also coupled in (1) to the Q point mass, but are not represented in the Figure.

In Fig. 1 what we see is that Q oscillates harmonically for a couple of periods, indicating that the point mass is moving in a quadratic potential. This point mass is also subject to other forces, however, and is not simply a harmonic oscillator, because after a few oscillations, it suddenly begins a 'driving phase' of motion at nearly steady velocity, as if an engine had been turned on. After some time of this nearly steady motion, Q abruptly resumes harmonic motion, and at a much larger amplitude than it had initially; something has done a lot of work on the mass during the driving phase of motion. In fact this work has come from the system's other two oscillators, which evolve on high frequencies. Small effects of these coupled high-frequencies are shown, in the insets of

FIg. 1b), in the otherwise steadier speed of Q. The remarkable feature of Fig. 1, however, is a much more dramatic effect of the coupling of Q to the high-frequency oscillators. During the driving phase of motion, when $\dot{Q}$ stays nearly constant, the high-frequency part of the system steadily transfers energy to Q.

Figure 1 is thus a proof by example that daemons are allowed by the laws of mechanics [8]. We will further show that daemons are a general class of dynamical systems featuring three distinct time scales, and that they obey certain non-trivial constraints, which are in some respects reminiscent of the laws of thermodynamics. We will approach these general points by first explaining the particular system of Fig. 1.

II. A HAMILTONIAN DAEMON

A. The Hamiltonian

The Hamiltonian for the system of Fig. 1 is

$$H_1 = \frac{P^2}{2M} + \frac{M\nu^2}{2}Q^2 + \frac{p_+^2 + p_-^2}{2m} + \frac{m}{2}(\Omega_+^2 q_+^2 + \Omega_-^2 q_-^2) \\ -Kq_+q_- \cos(kQ)\ , \tag{1}$$

with the particular values $M = 500$, $\nu = 1$, $m = 1$, $\Omega_+ = 1500$, $\Omega_- = 500$, $K = 50\sqrt{3}$, $k = 100$. Since H_1 does not depend explicitly on time, it is conserved. The work done on the mass depicted in Fig. 1 is drawn from the initial energy of the Ω_+ oscillator, though the high frequency which actually matters dynamically in this system is really the difference $\Omega \equiv \Omega_+ - \Omega_- = 1000$. The quantitative parameters of this model were chosen for illustrative purposes; the system's behavior is not qualitatively sensitive to their precise values [8].

Secular energy transfer in H_1 is non-trivial: if we let $K \to 0$ we would simply have three harmonic oscillators, and even with the nonzero value of K that we actually have, their three frequencies all differ by much more than the frequency scale $\sqrt{K/m}$ associated with the coupling between them. From this one would normally deduce *adiabatic decoupling*: to zeroth order in K the three coordinates should simply oscillate, and do so at greatly different frequencies, so that the coupling term would be not only small, but rapidly oscillating, and hence it should average to zero and have no significant effect of any kind. Beyond zeroth order, energy exchange would indeed occur, but it would be small in amplitude and high in frequency, with no net long-term (secular) effect. This kind of adiabatically decoupled behavior is indeed seen in Fig. 1 before and after the driving phase, but the driving phase defies the expectation of adiabatic decoupling.

The expectation of adiabatic decoupling in this case is based on the assumption that a weak coupling, such as we have, can only have substantial effects if it is also resonant. In linear dynamics, resonance means frequency matching, and since our three oscillators' frequencies are badly mismatched, linear dynamics would imply adiabatic decoupling. Weak but highly nonlinear couplings, however, may be resonant after all, at least in some regions of phase space, even across large frequency gaps. This is the kind of exception to adiabatic decoupling that our daemon exploits.

B. Adiabatic decoupling vs. nonlinear resonance

Numerical solution of the equations of motion for H_1 in fact confirms that adiabatic decoupling is the case throughout most of phase space. The coupling term in H_1 is a highly nonlinear function of Q, however, and this provides a nonlinear loophole in adiabatic decoupling. A special case occurs when $Q(t)$ happens to be rising or falling at close to the critical speed $v_c = \Omega/k = 10$. In this circumstance, $q_+q_- \cos(kQ)$ is not just rapidly oscillatory, so as to average to zero, but includes a component which changes slowly. Whenever $\dot{Q} \sim \pm v_c$, therefore, the coupling may indeed have a more significant effect, at least for the short time during which the speed remains close to v_c[9].

An effect that is only significant for a short time is not really very significant, however; it would not count as an exception to adiabatic decoupling, which is about long-term behavior. But this is our key point, because as we will see below, when $\dot{Q}$ is near v_c, the significant effect of the coupling *can be to keep* $\dot{Q}$ *near* v_c. So instead of a transient perturbation, the nonlinear resonance at $\dot{Q} = v_c$ can induce a transition to a new dynamical phase — the driving phase. As Fig. 1 shows, $\dot{Q}$ can remain close to v_c for a long time, until the slow mass has gained a large amount of potential energy from its steady climb up the harmonic potential $M\nu^2Q^2/2$.

C. Canonical transformation to H_2

To see how the daemon works, we will apply canonical transformations and adiabatic approximations to transform H_1 into a series of effectively equivalent H_n. The first step is a canonical transformation to Hamilton-Jacobi variables (τ, U), (α, A) for the two fast oscillators:

$$q_\pm = \frac{\sqrt{\pm 2(U - \Omega_\mp A)}}{\sqrt{m\Omega\Omega_\pm}} \cos(\Omega_\pm \tau + \alpha) \\ p_\pm = -\frac{\sqrt{\pm 2m(U - \Omega_\mp A)\Omega_\pm}}{\sqrt{\Omega}} \sin(\Omega_\pm \tau + \alpha)\ . \tag{2}$$

If we then discard small terms that play no significant role in the Fig. 1 evolution because they are never resonant for $\dot{Q}$ near $+v_c$ (see our Appendix for details), we obtain the adiabatic effective Hamiltonian

$$H_2 = \frac{P^2}{2M} + \frac{M\nu^2}{2}Q^2 + U \tag{3}$$

$$-\kappa\sqrt{(\Omega_+ A - U)(U - \Omega_- A)}\cos[k(Q - v_c\tau)]$$

$$\kappa = \frac{K}{2m\Omega\sqrt{\Omega_+\Omega_-}} = 5\times 10^{-5} \text{ in Fig. 1 .}$$

Since the angle variable α does not appear in H_2, its conjugate momentum A is a constant of the motion under the adiabatically approximate evolution generated by H_2 (hence an adiabatic invariant in the exact H_1 evolution). We may therefore consider A as a constant, whose value is given by $A(0)$, in the Hamiltonian evolution of Q, P, τ and U. Inverting (2) yields $A(0) = 8\times 10^4$ in the case of Fig. 1.

Comparison with a Chirikov resonance

(Readers who are not interested in the relationship between our model and an isolated Chirikov resonance may skip directly to the next subsection.)

The coupling term in H_2 of (3) is now similar in mathematical form to the simplest case (an isolated resonance) of the resonances studied by Chirikov in [7]. One apparently significant difference, however, is that our τ is a canonical variable, which evolves under the time-independent Hamiltonian of the total system, whereas the corresponding variable (also called τ) in Ref. [7] is simply a multiple of the time, and appears as the argument of a sinusoidally time-dependent external perturbation. This difference is not really as sharply qualitative as it might seem, though, because as Chirikov points out, a fixed time dependence of the form $\tau \propto t$ can always be regarded as the evolution of a dynamical variable like our own τ, if a term linear in τ's conjugate variable U is added to the Hamiltonian. As long as U appears nowhere else in the Hamiltonian, then, the equation of motion for τ will simply be $\dot\tau = 1$.

With this understanding (and after minor adaptations of notation) we can therefore express the Chirikov Hamiltonian of Eqn. (3.15) in Ref. [7] as

$$H_{Ch} = K(I) + U + G(I)\cos(k\theta - \Omega\tau) \,, \qquad (4)$$

where θ, I and τ, U are canonically conjugate pairs. Thus the first difference between our model and Chirikov's is that the factor which multiplies our cosine coupling depends on U instead of I. We will see in Section IV, below, that this dependence on U is important in ensuring that the daemon's downconversion must eventually cease if U is bounded from below (*i.e.* our fuel source is finite). For the ongoing operation of the downconversion mechanism, however, this detail of coupling dependence on U is not really significant, because our U will only change slowly.

The most basic difference between our H_2 and the above H_{Ch}, in fact, is that in H_{Ch} the slow variable θ which is coupled to the fast τ does not appear in $K(I)$, which depends only on the conjugate variable I. In our case, in contrast, the variable coupled to τ is Q, and the slow mass's own Hamiltonian includes not only a term that depends on its conjugate, P (namely the kinetic energy $K(P) = P^2/(2M)$), but also a term that depends on Q itself (the slow-sector potential energy $V_S(Q) = M\nu^2Q^2/2$). As we will explain just below, the precise form of this potential term in our model is not important for the downconversion mechanism; all that really matters is that we have a sufficiently slowly varying potential $V_S(Q)$. In this sense the Chirikov resonance is a degenerate case of our daemon model, in which $V_S(Q)$ happens to vanish. As we will see further on in this Section, the existence of a non-trivial potential term makes a qualitative change in our model's dynamics, in comparison with that of the isolated Chirikov resonance which our model otherwise resembles. The additional potential term introduces an additional time scale, and the phenomena that occur on this long additional time scale are the main subject of this paper: the long additional time scale corresponds precisely to the rate of steady downconversion. In the Chirikov limit H_{Ch}, this rate goes to zero.

D. Slowly varying potential

The harmonic potential $M\nu^2Q^2/2$ changes only very slightly over the period of the rapid oscillations in $\dot Q(t)$ that occur during the driving phase. We can therefore understand the dynamics in this phase by replacing the harmonic potential with a linear potential whose slope g is constant, and then after this simplified problem has been solved, replacing g with the slowly changing slope $M\nu^2Q(t)$ of the actual harmonic potential. This yields

$$H_3 = \frac{P^2}{2M} + gQ + U \qquad (5)$$
$$-\kappa\sqrt{(\Omega_+ A - U)(U - \Omega_- A)}\cos[k(Q - v_c\tau)] \,.$$

In fact we could apply a similar procedure to replace any sufficiently slowly varying $V_S(Q)$, not necessarily quadratic, with an instantaneous linear potential whose slope changes slowly in comparison with the time scale that is relevant for the downconversion mechanism. We have confirmed with numerical solutions of such nonlinearly modified alternatives to H_1 that this procedure accurately approximates the exact evolution.

E. An explicitly time-dependent invariant

We now reach a crucial point in daemon dynamics. Linearizing the harmonic potential of H_2 has given the new Hamiltonian H_3 a first integral:

$$J = U + (P + gt)v_c \qquad (6)$$

is constant (even though it depends explicitly on t!) under the evolution generated by H_3. In the adiabatically

decoupled phase, J stays constant by having $\dot{U} \doteq 0$ and $\dot{P} \doteq -g$; in the driving phase it is constant because $\dot{U} \doteq -gv_c$ and $\dot{P} \doteq 0$. For the evolution shown in Fig. 1, $J \doteq 1.2 \times 10^8$.

Using the first integral J we can reduce H_3 to a single non-cyclic degree of freedom by a further canonical transformation from $(Q,P),(\tau,U)$ to $(q,p),(\tau,J)$, with $q = Q - v_c\tau$ and $p = P - Mv_c$. With the explicit t-dependence of J, this is a time-dependent canonical transformation, and the transformed Hamiltonian which results [10] is

$$H_4 = \frac{p^2}{2M} + gq - \kappa\sqrt{(\Omega_+ A - \tilde{U})(\tilde{U} - \Omega_- A)}\cos(kq)$$
$$\tilde{U}(p,t) = J - (p + Mv_c + gt)v_c \; . \qquad (7)$$

Since J is constant, we can now treat H_4 as a time-dependent Hamiltonian in the two-dimensional phase space (q,p). Note in comparison that a pure Chirikov resonance, with $g \to 0$ as in H_{Ch} above, would have a time-independent Hamiltonian in place of H_4, as in the 'resonance Hamiltonian' H_r of Eqn. (3.18) in Ref. [7].

F. A 'tilted washboard' model

Finally, the p-dependence of $\tilde{U}$ can be important during the dynamical transition from decoupling to driving, but during bound motion the kinetic term in H_4 must necessarily be smaller than the depth of the local potential minimum, and so $p \lesssim \sqrt{M\kappa J}$. This means that the proportional contribution of pv_c within $\tilde{U}$ is $\mathcal{O}(\sqrt{\kappa M v_c^2/J}) \sim 10^{-5}$. We can therefore understand both the driving and decoupled phases by neglecting pv_c in $\tilde{U}$, obtaining for our final approximate adiabatic Hamiltonian simply that of a single particle in a time-dependent potential:

$$H_5 = \frac{p^2}{2M} + V(q,t)$$
$$V(q,t) = gq - \kappa\sqrt{u(t)[\Omega A - u(t)]}\cos(kq)$$
$$u(t) = J - \Omega_- A - Mv_c^2 - gv_c t \equiv u_0(1-\gamma t) \; . \quad (8)$$

We emphasize that the time-dependent effective potential $V(q,t)$ which appears in the final approximate H_5 is not the same as the original slowly varying external potential $V_S(Q) = M\nu^2Q^2/2$ from the full and exact H_1. See our Appendix for numerical confirmation that H_5 accurately reproduces the exact evolution under H_1, as shown in Fig. 1, except for oscillations on the short $1/\Omega$ time scale.

Inserting the adiabatic time dependence $g = M\nu^2 Q(t) \doteq 500 \times 10(t - 12.5)$ reveals that the rate γ in $u(t)$ as just defined is of order 10^{-2} throughout the evolution shown in Fig. 1. This is by far the lowest rate or frequency in the entire evolution. The effective potential $V(q,t)$ can thus be considered adiabatically, by understanding the motion implied by $V(q,t)$ for fixed t, and then allowing the parameters in that motion to change slowly with t. For any fixed t, V has the form of a tilted sinusoidal 'washboard'; see Fig. 2. The two dynamical phases of H_1 are then finally clear: for most q,p orbits the slight undulations of the potential are insignificant (adiabatic decoupling), but as long as the washboard tilt is not too great, there exist local minima in q of $V(q,t)$, about which bound motion is possible in the surrounding local potential well.

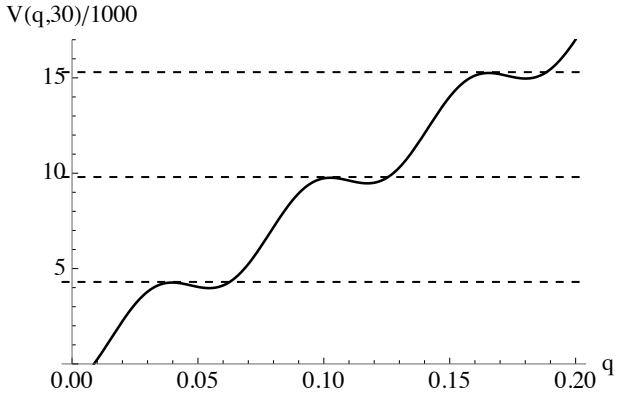

FIG. 2. Effective potential $V(q,t)$ from H_5, which represents H_1 in new canonical variables, with linearized external potential and adiabatic elimination of non-resonant terms. The time $t = 30$ is chosen for illustration; the potential looks similar for other t during the driving phase of Fig. 1. Dashed lines are horizontal guides to the eye, to confirm that shallow wells exist. Unbound orbits correspond to the adiabatically decoupled phase of H_1; bound orbits, to the steady downconversion of the daemon's driving phase.

When p thus oscillates around zero, P oscillates slightly around Mv_c, and hence $\dot{Q}$ oscillates slightly around the steady speed v_c, as seen in Fig. 1 for $13 \lesssim t \lesssim 51$ (the driving phase). (These small oscillations of q,p in the local well of V are the intermediate-scale oscillations seen in the upper inset of Fig. 1b).) The crucial point is that by maintaining the mass's speed near v_c, the rest of the system steadily does work on the mass, by raising it steadily against the gQ or $M\nu^2Q^2/2$ potential. This is how the daemon performs steady downconversion.

G. Self-sustaining resonance

The simple picture of effective particle motion that we finally reached in $H_5(q,p,t)$ does let us understand the daemon's driving phase in familiar terms, but only through the rather abstract variables that emerged after a series of transformations and adiabatic approximations. The Hamiltonian daemon represented by H_1 is not a trivial system.

One somewhat more physical way to understand the real nature of the steady downconversion mechanism is to consider it as a nonlinear version of the ordinary kind

of resonant energy transfer, with the special nonlinear feature being that the resonance only holds for $\dot{Q}$ close to v_c. As usual, therefore, energy transfer will be most efficient at exact resonance, which occurs when $\dot{Q}$ is exactly v_c. The mass might therefore rise stably at a speed just above v_c, so that if it were to decelerate, transfer efficiency would rise and the additional force would speed the mass back up. Conversely, if it were to accelerate, transfer efficiency would fall, and the mass would slow down. The nonlinear resonance can therefore be self-sustaining.

This effect is an example of what is generally known as *nonlinear stabilization* [7]; but the effect of the stabilization is here somewhat different from usual. Ordinarily what is meant by nonlinear stabilization is that nonlinear effects *prevent* the secular transfer of energy between degrees of freedom. In a Chirikov resonance as described by H_{Ch} in (4) above, for example, a linear $K(I) = \omega I$ would mean that resonance ($k\omega = \Omega$) would obtain, or not obtain, throughout the whole phase space uniformly, depending only on the fixed parameters in the Hamiltonian (namely $k\omega$ and Ω). If resonance does hold, therefore, the special case of linear $K(I)$ allows complete transfer of energy from U to I. As well as being special, the case of linear $K(I)$ is only a degenerate case of downconversion, because the designation of τ as fast and θ as slow is in this case arbitrary: one could give the two oscillators the same frequency by the canonical rescaling $\theta = X\tilde{\theta}$, $I = \tilde{I}/X$ for $X = \omega/\Omega$.

If $K(I)$ in H_{Ch} is nonlinear, however, then the resonance condition $kK'(I) = \Omega$ will generally only hold for a particular value of I — just as in our case the quadratic kinetic energy $P^2/(2M)$ means that resonance occurs only at $P = Mv_c$. If I is then increased in the Chirikov model, due to resonant transfer from U, resonance no longer holds, and transfer stops. The result in the Chirikov model is that I can then only oscillate around the resonant value, with an amplitude limited by the strength of the cosine perturbation amplitude $G(I)$; the nonlinear stabilization prevents the secular transfer of energy from U to I. In our case also the resonant transfer phase of evolution (the driving phase) sees only small oscillations of P around the resonant value Mv_c, and the amplitude of these oscillations is limited by our interaction strength κ. Since our slow mass's energy also includes a potential term $V_S(Q) \to gQ$, however, the nearly steady maintenance of $P \doteq Mv_c$ means that the potential energy steadily increases, through resonant transfer from U. In our daemon model, therefore, the nonlinear stabilization of P therefore means precisely that secular energy transfer is not prevented, but sustained.

This is the basic mechanism of steady downconversion in Hamiltonian daemons like the one described by our H_1. The important consequence is that total energy transfer from U to the slow mass is not limited by the strength of the coupling term, but by the total amount of available U (and by other constraints due to the slow γ-time-scale time dependence of H_4, to be described below). The strength of the coupling term does limit the daemon's performance; but it limits it rather in the way that a car engine's performance is limited by its maximum torque output. A car whose engine is too weak cannot climb too steep a hill. If the torque suffices for the hill's steepness, however — corresponding here to the tilt of $V(q,t)$ being gentle enough, in relation to the well depth, that local minima still exist — then the car can keep driving uphill until it runs out of fuel.

III. GENERIC DAEMONS

The (mathematical) existence of Hamiltonian H_1 proves that Hamiltonian daemons exist. How representative is H_1 of daemons in general? In a word, broadly. It is representative of all daemons whose slow 'work' degree of freedom is powered from high-frequency 'fuel' that is not in itself chaotic, and which exploit isolated nonlinear resonances. For all such systems, the high-frequency sector can be represented without loss of generality with Hamilton-Jacobi co-ordinates τ, U, plus as many additional co-ordinates $\vec{\alpha}$ as may be needed for all the high-frequency degrees of freedom. Since the Hamilton-Jacobi construction makes the $\vec{\alpha}$ time-independent under the evolution generated by U alone, the $\vec{\alpha}$ can only evolve under the Hamiltonian terms which couple U to the slow degree of freedom. These terms have to be small, because we can only identify the slow and fast degrees of freedom as such by overlooking their coupling on the grounds that it is weak. The $\vec{\alpha}$ variables will therefore evolve slowly under the full daemon H, and so considerations of resonance and adiabatic invariance like those applied to H_1 will also lead, for a large class of daemon systems, to explicitly time-dependent adiabatic invariants like J, and thus to time-dependent effective Hamiltonians of the forms of H_4 and H_5. The potential function V may be different in its dependence on u as well as on q, and it may have further slow time dependence through its dependence on $\vec{\alpha}$ variables. The basic scenario, of steady downconversion maintained by cyclic orbits in an effective potential, will remain the same.

In particular it is interesting to note that daemons are generically dynamical systems with at least three important time scales. By their definition they have an important short time scale, analogous to Ω^{-1} in H_1. Also by the definition of steady downconversion, they must have a much longer time scale, analogous to γ^{-1} in H_5, over which the high-frequency energy is transferred into slower motion. The steady slowness of this transfer is ultimately also due to the smallness of the interaction between the slow and fast subsystems; if it were not so slow, then the energy transfer would not really be analogous to the steady downconversion of an engine, but to

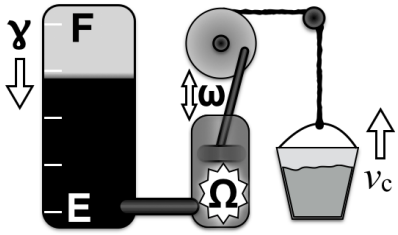

FIG. 3. Schematic: Daemons as systems with three time scales, analogous to the dynamical time scales of a combustion engine. Energy is stored densely in degrees of freedom with a high dynamical frequency Ω. This energy is consumed at a slow rate γ to do work at a steady speed v_c. The fact that $\gamma \ll \Omega$ distinguishes engine-like dynamics (steady downconversion) from bomb-like (rapid energy release). The non-trivial mechanism which achieves steady downconversion involves cycles on the intermediate frequency scale ω, with $\gamma \ll \omega \ll \Omega$.

the explosively non-steady downconversion of a bomb.

It is the intermediate third time scale which is surprising here: the period of effectively bound motion during the driving phase, as shown in the upper inset of Fig. 1b. The generic requirement for (adiabatically) periodic motion of this kind is clear from the nature of dynamical stabilization: Hamiltonian systems maintain a steady average speed by oscillating around it. If we label the oscillation frequency with ω, then the weakness of coupling must imply $\omega \ll \Omega$, while the steadiness of downconversion implies $\omega \gg \gamma$. Daemons are thus examples of multiple-time-scale systems in which three distinct time scales are simultaneously important. Whether by coincidence or not, the essential feature of steady periodic motion in a daemon's driving phase is analogous to the cyclic operation of a macroscopic engine. We represent this fact provocatively in Fig. 3 by comparing Ω, ω, and γ to the analogous rate scales of a motor.

Daemons are thus a large class of systems, but very many of them can be mapped onto systems similar to the one presented here. The exploration of possible exceptions and generalizations must be left for future work. Here we will turn instead to examining an important general feature of daemons which can also be understood from the behavior of H_1. How and why do daemon driving phases end?

I.: EFFICIENCY LIMITS

A. Energy conservation

The explicit slow γt time dependence of H_5 is *not* any kind of heuristic modification of the system dynamics, but is due entirely to the time-dependent transformation which replaced U with J. That step, which produced H_4 from the time-independent H_3, was an exact canonical transformation. Since as a result H_5 is explicitly time-dependent, it is not conserved; but the original energy H_1 still is conserved. Indeed H_5 is time-dependent precisely *because* H_1 is conserved, since the explicit time dependence of the adiabatic invariant $J = U + (P + gt)v_c$ reflects the fact that if $P \doteq Mv_c$ remains (nearly) constant despite the external force $-g$, then U must decrease steadily to supply the needed power gv_c. It should therefore be no surprise that it is the slow time dependence of $V(q,t)$ in H_5 which makes the daemon stop driving when it runs out of U 'fuel'. And analogous phenomena will be observed in any analogous daemon system.

In fact the H_5 daemon must in general stop *before* it consumes all its U. Whether because the washboard tilt g increases, or because the washboard well depth $\kappa\sqrt{u(\Omega A - u)}$ eventually decreases, at some time t_0 there will no longer be any local minima in V. No bound orbits, and hence no steady downconversion, can continue past t_0, even if $U(t_0)$ has not yet been drained to its ground state. At the minimum value $U = \Omega_- A$, the well depth is zero; but with finite g, V will become a bumpy slope with no minima before the washboard completely disappears. A similar consideration will apply to any daemon whose effective Hamiltonian is adiabatically similar to H_5.

B. Efficiency limits from adiabaticity

Except for a zero-measure set of initial conditions, moreover, daemons must stop driving, and revert to adiabatic decoupling, even earlier than their t_0. This more stringent limit on Hamiltonian work extraction from high-frequency energy is the subtle inherent weakness in steady downconversion. The daemon can run for a long time, because it drains its fuel slowly: $\gamma \ll \omega$ is slow. But this means that, in the driving phase, the area S that is enclosed in (q,p) phase space by the periodic orbits of bound motion is an adiabatic invariant [10]. At any $t < t_0$, any local well of V will have a highest bound orbit, and the area $S_{\max}(t)$ which it encloses in the (q,p) plane will be the largest S that the local well can support at that time. For any steady downconversion trajectory, therefore, the orbit's enclosed area S will remain constant (adiabatically, *i.e.* up to small, fast oscillations), while the maximum area $S_{\max}$ which the well can support will eventually fall towards zero. After the time $t_S < t_0$, when $S_{\max}(t_S) = S$ has fallen to S, the bound orbit with area S can no longer continue, and the dynamical transition to adiabatic decoupling must occur. See Fig. 4.

The adiabatic invariance of S thus sets a limit on how much high-frequency energy a daemon can convert into work, which is in general lower (and as Fig. 4 shows, possibly much lower) than the limits set by energy conservation, or by the maximum force that can be exerted by the coupling term in H_1. This limitation is moreover inevitable, because it is inherent in the time scale

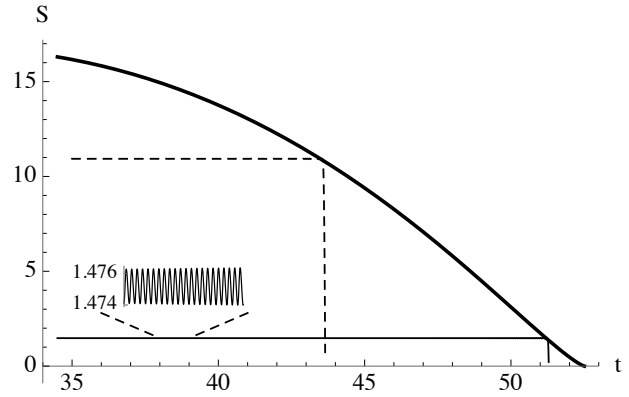

FIG. 4. Area $S(t)$ enclosed in (q, p) phase space by bound orbits under H_5, for the parameter values of Fig. 1.
Thick curve: S for the largest bound orbit that is possible at time t. This maximum area shrinks as the lattice tilt increases and the local minima become shallower. At time $t_F = 52$ the harmonic force $-g(t)$ exceeds the maximum force the lattice can exert, and no minima exist.
Solid curve: S for the trajectory shown in Fig. 1. Since the explicit time dependence of H_5 is slow, S is adiabatically invariant; when the potential minimum becomes too shallow to contain it, the dynamical transition to adiabatic decoupling must occur.
Dashed curve: S for a less efficient trajectory, under H_1 but with different initial conditions. Here the adiabatic invariance of S forces the daemon to stop even though further driving is not forbidden by force limits or energy conservation.
Inset: The adiabatic invariant S is not strictly constant, but exhibits small, rapid oscillations around an average value that holds steady over long times.

hierarchy that defines a daemon, as sketched in Fig. 3. This generic daemon time scale hierarchy means that S must be adiabatically invariant, and limit the duration of steady downconversion to be less than what would be allowed by conservation of energy alone.

V. OUTLOOK TO MICROTHERMODYNAMICS

The main result of this paper has been to prove by example that Hamiltonian daemons are in principle possible; small dynamical systems can sustain steady energy transfer from rapid motions to slow. We have also found a non-trivial limit on the amount of work that our example daemon can do, and as we have argued that our example daemon is a representative example of a large class of daemon systems, we have shown that this kind of efficiency limit is a generic feature of steady downconversion in small Hamiltonian systems. Small dynamical systems have traditionally been studied as toy models for more general phenomena, however. Since the behavior of daemons is at this level analogous to that of combustion engines, and since heat is after all fundamentally nothing but a certain rapid 'kind of motion' [13], we close our paper by discussing the possible implications of our results for the microscopic roots of thermodynamics.

A. Daemons and thermodynamics

The fact that Hamiltonian daemons can in general do less work than the limit allowed by energy conservation, because a certain phase space volume cannot decrease, is reminiscent of the thermodynamic limits on heat engine efficiency, which not only insist on conservation of total energy, but further say that not all energy is thermodynamically 'available' for conversion into work, because entropy cannot decrease. In the statistical mechanics of the microcanonical ensemble, entropy is identified with (the logarithm of) a certain phase space volume (namely, that volume which is ergodically explored by the system as it undergoes chaotic evolution). It is therefore intriguing to speculate whether the mechanical constraints on Hamiltonian daemons, in which conservation of certain phase space volumes limits energy transfer, may perhaps be related somehow to the constraints of thermodynamics, whose microscopic origins remain obscure.

Further investigations can moreover demonstrate additional analogies between daemon mechanics and thermodynamics, in the constraints on daemon ignition, and in entanglement growth in quantum mechanical daemons; these results are reported elsewhere [11, 12]. These analogies may even prompt an alternative hypothesis to the long-standing assumption that thermodynamics somehow 'emerges', at mesoscopic scales, from non-thermodynamical microphysics. The alternative is that perhaps thermodynamics doesn't begin at mesoscopic scales, but simply persists into macroscopic scales from microscopic mechanics, where it already exists in embryonic form as the set of dynamical constraints upon steady downconversion.

On the other hand we must note that all the constraints on Hamiltonian daemons obtain within closed-system Hamiltonian mechanics, even without chaos or disorder. While the constraints on daemons may therefore resemble thermodynamics in their practical consequences, the reasons why the constraints apply would seem to be different. For the constraint that we have identified in this paper, for example, the difference between statistical mechanical entropy and the adiabatic invariant S is drastic in a basic respect: the phase space area involved in limiting daemons does *not* represent a volume that is ergodically explored by the system. The system's orbit under H_5 *encloses* the area S, but it does not even enter the area S, let alone ergodically fill it. The concept of ergodic filling is fundamental in statistical mechanics, and if ergodic filling is really irrelevant to the mechanical constraints on steady downconversion, then the relationship between limits on daemons and thermodynamics may go beyond statistical mechanics itself,

and lead either to an extension of thermodynamics into smaller systems ('microthermodynamics'), or the embedding of thermodynamics within a larger class of analogous constraints.

B. Daemons and baths

Should we really be certain, in any case, that statistical mechanics is the whole story of how thermodynamics emerges from mechanics? Statistical mechanics is based on introducing ensembles to provide 'coarse-grained' descriptions, which contain no information about any aspects of time evolution that are rapidly ergodic. In this sense the first question that statistical mechanics addresses is about the exchange of *information* between fast and slow variables, and its first answer is that equilibration occurs: information can be effectively lost from the slow sector, even though the total evolution is Liouvillian (or in quantum mechanics, unitary). Questions about energy exchange are then most often treated in a second step, in the sense that one studies energy transport under local (but not global) equilibrium, or postulates 'baths' or 'reservoirs' as sinks or sources of energy that always remain in equilibrium.

Active research continues into how information can thus effectively be lost within deterministic time evolution, and the physics of information is a major topic in general. By no means do we question the profound importance of information in physics; but we would argue also for the alternative approach, of considering energy exchange as the first question, instead of equilibration. Thermodynamics is after all a practical discipline, and engines do not exist to equilibrate, but to perform large amounts of steady work from compact fuel.

There is a basic physical reason why extracting large amounts of work from compact fuel demands downconversion: the units of energy are Mass $\times$ (Length)^2 / (Time)^2, and so a source of energy which is rich in comparison to its size and mass must possess a short time scale. It may therefore be worthwhile, as a kind of second front in the struggle to understand the microscopic roots of thermodynamics, to re-phrase the basic question. Instead of asking *How can steady work be extracted from the equilibrated degrees of freedom of a heat bath?*, we can ask *How can steady work be extracted from high-frequency degrees of freedom?* Further investigations of Hamiltonian daemons and related systems may thereby shed further light on thermodynamics.

Understanding Hamiltonian daemons may potentially even be practical. Chemical reactions are a ubiquitous source of this kind of dense energy, and the general problem of transferring their high-frequency energy into much slower motions is one that is not only solved macroscopically by engines, but also (somehow) by living organisms, on much smaller scales. Whether dynamical limitations on steady downconversion in small systems prove to be a microscopic limit of thermodynamics, or whether they turn out to be a parallel set of constraints, they may conceivably be of comparable importance to thermodynamics, for molecular machines [14].

Of course natural molecular machines are not isolated systems as daemons are, and any future artificial microscopic machines will also have to tolerate dissipation and noise if they are to be useful. We have only defined daemons as isolated, however, in order to ensure that they do not *need* external power in order to operate. The adiabatic mechanisms by which daemons operate are normally quite robust, and so a generalized daemon, which was subject to dissipation and noise from its environment, might still very well function. If so, the processes of self-contained steady downconversion as by daemons, and work extraction from baths as by heat engines, may prove to be alternatives that are compatible — or even to be alternative limits of a single phenomenon.

[1] M. Born and V. Fock, Beweis des Adiabatensatzes *Z. Phys. A* **51**, 165–80 (1928)

[2] M. Born and R. Oppenheimer, Zur Quantentheorie der Molekeln *Ann. Phys.* **389** 457–84 (1927)

[3] S. Weinberg, *The Quantum Theory of Fields: Volume 2, Modern Applications*, Cambridge University Press (2005)

[4] R. Shankar, Renormalization-group approach to interacting fermions *Rev. Mod. Phys.* **66**, 129 (1994)

[5] J. Berges, N. Tetradis, and C. Wetterich, Non-perturbative renormalization flow in quantum field theory and statistical physics *Phys. Rep. 363*, 223–386 (2002)

[6] V.I. Arnold, *Mathematical methods of classical mechanics*, Springer (1989)

[7] Boris V. Chirikov, A universal instability of many-dimensional oscillator systems. *Phys. Rep.* **52**, 263 (1979)

[8] At least 25 digits of numerical working precision are required to reproduce the plotted evolution of Fig. 1 from the stated initial conditions. The highly nonlinear dynamics involves non-adiabatic crossing of a separatrix, and the precise evolution is therefore quite sensitive. Numerically experimenting with different conditions readily confirms, however, that fine-tuned initial conditions are *not* required in order to obtain qualitatively similar driving behavior. High precision is merely needed to obtain a precise reproduction of the particular evolution shown, which was chosen to exhibit all the significant features clearly and within a convenient time frame.

[9] This explains the slight anharmonic perturbations in $\dot{Q}(t)$ which can be seen in the first 10 time units in Fig. 1, whenever $\dot{Q} \sim \pm 10$.

[10] H.P. Goldstein, P. Charles, and J.L. Safko, *Classical mechanics*, Addison Wesley (2002)

[11] L. Gilz, E.P. Thesing and J.R. Anglin, http://arxiv.org/abs/1509.02778

[12] L. Gilz, E.P. Thesing and J.R. Anglin, http://arxiv.org/abs/1304.3222

[13] R. Clausius, Über die Art der Bewegung, welche wir

Wärme nennen, *Annalen der Physik* **176**, 353–380 (1857)
[14] R.D. Vale and R.A. Milligan, The Way Things Move: Looking Under the Hood of Molecular Motor Proteins, Science **288**, 5463 (2000)

ACKNOWLEDGEMENTS

We thank our colleague Martin Strzys for suggesting 'microthermodynamics' as a name for the hypothesis that thermodynamics extends, in some non-trivial form, to small systems. LG acknowledges funding from the German Excellence Initiative (DFG/GSC 266).

APPENDIX: ACCURACY OF H_5 AS AN ADIABATIC APPROXIMATION TO H_1

Non-resonant terms discarded to obtain H_2 from H_1

A crucial step in our main text's chain of effective Hamiltonians $H_{1\to5}$ is discarding non-resonant terms from H_2. Without discarding them, the canonical transformation to Hamilton-Jacobi variables would instead have produced from H_1 the canonically equivalent Hamiltonian

$$\begin{aligned} H_{1.5} &= \frac{P^2}{2M} + \frac{M\nu^2}{2}Q^2 + U - \kappa\sqrt{(\Omega_+ A - U)(U - \Omega_- A)}\tilde{K}(\tau,\alpha,Q) \\ \tilde{K} &\equiv 4\cos(kQ)\cos(\Omega_+\tau+\alpha)\cos(\Omega_-\tau+\alpha) \equiv 2\cos(kQ)\Big(\cos[(\Omega_+ + \Omega_-)\tau + 2\alpha] + \cos(\Omega\tau)\Big) \\ &\equiv \cos[kQ + (\Omega_+ + \Omega_-)\tau + 2\alpha] + \cos[kQ - (\Omega_+ + \Omega_-)\tau - 2\alpha] + \cos(kQ + \Omega\tau) + \cos(kQ - \Omega\tau)\ . \end{aligned} \tag{9}$$

The four terms in $\tilde{K}$ will then be resonant for $\dot{Q} \doteq \mp(\Omega_+ + \Omega_-)/k$ (first two terms) and $\dot{Q} \doteq \mp\Omega/k$ (last two terms), respectively. Non-resonant terms can be neglected, because of adiabatic decoupling; this is essentially the so-called "rotating wave approximation" of quantum optics. Since the only non-trivial motion that occurs in the trajectory shown in Fig. 1 involves $\dot{Q} \doteq v_c = \Omega/k$, the first three terms can be discarded, as far as this motion is concerned. This leaves us with H_2 from Eqn. (3) of the main text. The accuracy of this rotating wave approximation, and of the other approximations which lead from H_1 to H_5, may then be assessed directly, by numerically solving the equations of motion for both Hamiltonians.

Accuracy of H_5 as an approximation to H_1

To compare evolution under the explicitly time-dependent adiabatic effective Hamiltonian H_5 to the full evolution under the original time-independent Hamiltonian H_1, we must relate the q, p coordinates to the original coordinates. Note that q and p are exactly canonical, even though non-resonant terms have been dropped between H_1 and H_5 so that H_5 is only approximately equivalent to H_1. The momentum transformation is just a translation: $p = P - Mv_c = P - 5000$. After inverting Eqn. (2) to obtain τ in terms of $q_\pm, p_\pm$ we find $q = Q - v_c\tau$ to be

$$q = -\frac{1}{k}\cos^{-1}\left(\frac{(\Omega_+\Omega_- q_+ q_- + p_+ p_-)\cos(kQ) - (\Omega_+ p_+ q_- - \Omega_- p_- q_+)\sin(kQ)}{\sqrt{(p_+^2 + \Omega_+^2 q_+^2)(p_-^2 + \Omega_-^2 q_-^2)}}\right)\ . \tag{10}$$

where $\cos^{-1}$ is defined to be positive. (It turns out that in the driving phase Q lags slightly behind $v_c\tau$, so that q is negative; hence the overall minus sign.) From the above coordinate mappings we can therefore compute the exact $q(t), p(t)$ from the numerical solution to the full equations of motion for all six original canonical coordinates under H_1, as plotted in Fig. 1, and compare this to adiabatically approximate evolution of q and p alone under H_5. We will refer to these solutions as (q_1, p_1) and (q_5, p_5), respectively.

To obtain (q_5, p_5), we must first determine H_5 explicitly by supplying the correct slowly-varying external force. The initial conditions of Fig. 1 provide

$$U(0) = \frac{m\Omega_+^2}{2}q_+^2(0) = 1.2\times10^8\ . \tag{11}$$

The daemon begins steady downconversion at around $t = 13$ — the transition is by definition not something that can occur at one precise instant, because it is a change from non-periodic to periodic motion in the (q, p) plane,

and the exact moment within the first period at which the motion first becomes periodic is not well defined. We can, however, extrapolate the nearly constant slope of $Q(t)$, during the driving phase, back to the time $t = 12.51$ at which $Q(t)$ would have been zero if the driving phase had begun from zero Q. We therefore deduce the adiabatically approximate

$$U(t) \doteq \tilde{U}(t) \equiv U(0) - \frac{M\nu^2}{2}[v_c(t - 12.51)]^2 \ . \quad (12)$$

This yields

$$H_5(q,p,t) = \frac{p^2}{2M} + M\nu^2 v_c(t - 12.51)q \\ -\kappa\sqrt{[\tilde{U}(t) - \Omega_- A][\Omega_+ A - \tilde{U}(t)]}\cos(kq) \ .$$

Adiabatic effective Hamiltonians like H_5 are in general determined by decomposing the exact evolution into a sum of slow and fast parts: $q(t) = q_s(t) + q_f(t)$, for example. Where adiabatic methods are warranted, this decomposition remains valid over long times, and the fast term q_f remains small in amplitude, as well as high in frequency, whereas the slow part q_s may slowly change by large amounts. The goal of adiabatic methods is to discard the fast parts like q_f and describe q_s accurately, over long times, with simpler equations; even with ample experience it is astonishing how well this can work. A subtlety which is sometimes overlooked in deriving adiabatic methods, however, is that the initial conditions which must be used for the adiabatic approximate evolution, in order to match the slow part of the exact motion, are generally not the same as the exact initial conditions, because the exact initial conditions represent the sum of the initial values of both fast and slow components. Slightly perturbed or renormalized initial conditions must therefore be used in the adiabatic approximation. Since the full and exact initial conditions determine the entire evolution, the correct adiabatic initial conditions can be derived as particular functions of the exact initial conditions; but it is usually easier simply to determine the adiabatic initial conditions by fitting the first few periods of adiabatic evolution to the exact evolution. The magic of adiabaticity is that this initially fitted approximation will then remain accurate for a long time.

For the evolution shown in Fig. 1 we solve the canonical equations of motion derived from H_5, from the initial time $t = 16$, shortly after the daemon's driving phase has begun. (The adiabatic method does not work as well right around the dynamical transition, where post-adiabatic effects are significant and other methods must be applied.) By tuning the initial values to $q_5(16) = -0.01$ and $p_5(16) = -100$, we can achieve the good agreement shown in Fig. A1 between H_5 evolution and the exact H_1 evolution with the initial conditions of Fig. 1 (which applied at $t = 0$), over the early time window $16 \leq t \leq 16.4$.

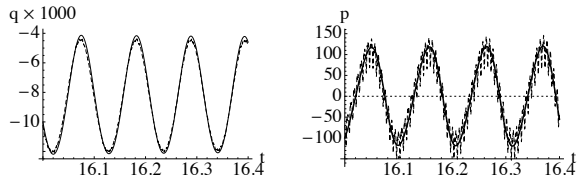

FIG. 5. Adiabatic evolution $(q_5(t), p_5(t))$ (solid curves) versus exact evolution $(q_1(t), p_1(t))$ (dashed curves), over the interval $16 \leq t \leq 16.4$. Initial conditions for the exact evolution are those stated in the caption of Fig. 1 in the main text, and apply at $t = 0$. Initial conditions for the adiabatic evolution are $q_5(16) = -0.01, p_5(16) = -100$. Fitting was done simply by eye. The suppression of high-frequency components from the adiabatic evolution is not error: it is what the adiabatic evolution is supposed to do.

We can then see how well this adiabatic solution q_5, p_5 tracks the exact evolution q_1, p_1 at later times. To do this, we are not allowed to retune the initial conditions for each later time period; we must keep the same initial conditions at $t = 16$, and see what they imply for q_5, p_5 later. Adiabatic evolution is not normally good, however, at following the exact phase of periodic motion over long times. Time translation is a zero mode, and the time coordinates of the exact and approximate solutions tend to drift slowly out of synchronization. We can therefore compare $(q_1(t), p_1(t))$ with $q_5(t - \Delta t), p_5(t - \Delta t))$, with Δt a new fitting parameter than is allowed to vary slowly in time. The agreement in phase that we then obtain, between exact and adiabatic solutions, is merely due to fitting; but the agreement in amplitude, and in all other slow features of the periodic motion, is non-trivial — indeed, remarkable. Over the long times that we investigate, the amplitude and frequency of both slow and fast components of the motion all change substantially, but the adiabatic approximation tracks the slow amplitudes and frequencies extremely well. See Fig. A2.

Adiabatic invariance of S under H_5 evolution

As well as being derived from H_1 as an adiabatic effective Hamiltonian, H_5 is itself a good target for adiabatic theory, because its explicit dependence on t is very slow compared to the frequency of the bound orbits it describes during steady downconversion. The slowness of the t-dependence of the functional form of H_5 naturally suggests an approximation based on the fact that H_5 remains *nearly* the same function over several periods. And indeed one can confirm that the instantaneous value $E_5(t) = H_5(t)$ for a particular trajectory changes only slightly, over several periods. The changes in $E_5(t)$ do not remain small over long times, however; $E_5(t)$ changes slowly, but it slowly changes by a large amount. When a Hamiltonian like H_5 is a slowly time-dependent function,

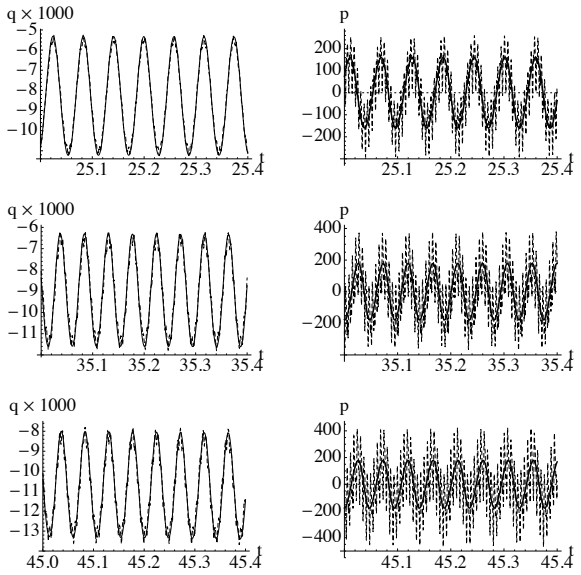

FIG. 6. Same as Fig. A1 above, but for time intervals $[25, 25.4], [35, 35.4]$ and $[45, 45.4]$, respectively, from top to bottom. Solid curves are the same adiabatic solution plotted in Fig. A1, but with fitted time shifts $\Delta t = -0.001$, $+0.025$, and -0.001, from top to bottom.

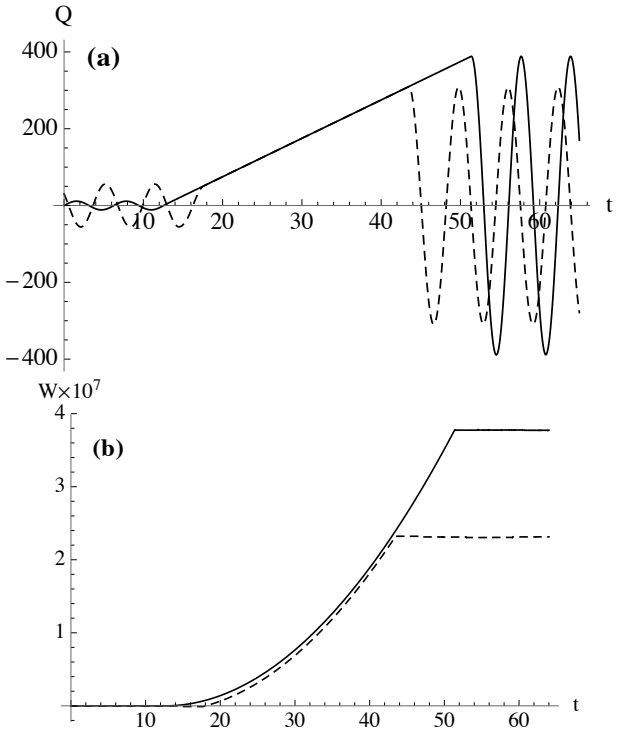

FIG. 7. **(a)** Same as Fig. 1a) of main text, but with the dashed trajectory from main text Fig. 4 for comparison. When both trajectories are in the driving phase, they are impossible to distinguish on this scale; this is a consequence of how the second trajectory was selected, as explained in the Appendix text. The dashed trajectory ceases steady downconversion at $t \doteq 43.7$, in agreement with Fig. 4.
(b) Work $W(t) = E_S(t) - E_S(0)$ done by the daemon on the slow mass, where $E_S = P^2/(2M) + (M\nu^2/2)Q^2$ is the energy of the slow oscillator. Final values $W(65)$ are 3.77×10^7 (solid) and 2.31×10^7 (dashed).

the adiabatic theorem states that the quantity which remains nearly constant over long times is not the value of the Hamiltonian, but rather the *action* $S(E_5, t)$.

The action S is defined for any fixed t and H_5 value E_5 to be the area in the (q, p) phase space plane which is enclosed by the closed curve $q_5(s), p_5(s)$ which satisfies $H_5(q_5, p_5, t) = E_5$. In our main text this action is computed and plotted in Fig. 4, for two different exact trajectories under H_1. For each trajectory the corresponding adiabatic $q_5(t), p_5(t)$ are determined by initial fitting, as described above. For each of the two trajectories shown in Fig. 4, the value $E_5(t) = H_5(t)$ is computed from the initially fitted $q_5(t), p_5(t)$, and then the corresponding enclosed phase space area S is computed by numerical integration. (No subsequently fitted time shifts Δt are employed in these computations, because $E_5(t)$ is a slowly varying function that is not sensitive to short time translations.) The plotted $S(t)$ are indeed close to being constant, until the transition to adiabatic decoupling occurs as explained in the main text.

As the inset in Fig. 4 shows, however, S is not an exact constant of the motion. As an adiabatic invariant, it exhibits small and fast oscillations around an average value that remains constant over a long time. For this reason the time at which steady downconversion ceases is not only intrinsically ill-defined within a period; it is impossible to predict from slow variables alone, because once the bound motion is only barely bound, a small post-adiabatic correction may be enough to let the system escape across the adiabatic H_5 separatrix into unbound motion. This adiabatic uncertainty in the precise time at which the daemon will stop working can easily be a very small fraction of the total duration of steady downconversion, however. For daemons whose time scale ratios Ω/ω and ω/γ are large, the significance of post-adiabatic corrections compared to adiabatic results may be comparable to the importance of thermal fluctuations relative to equilibrium thermodynamics in macroscopic engines.

The dashed line in the main text's Fig. 4 represents the same function S, computed for a slightly different trajectory $q_5'(t), p_5'(t)$, corresponding with the same adiabatic kind of accuracy to a different exact trajectory $q_1'(t), p_1'(t)$. To determine the particular different trajectory, a certain fine-tuning of initial conditions was employed for graphical reasons. The thick curve $S_{\max}(t)$ depends on J, and hence even though it is defined by the adiabatic Hamiltonian H_5, it really depends on the full initial conditions of the six canonical co-ordinates in the exact H_1. We were able to generate a second exact trajectory with (almost) exactly the same $S_{\max}(t)$ as

the trajectory from Fig. 1, so that it could be plotted in the same Fig. 4, by imposing initial conditions at time $t'_I = 35.0002$. The conditions we imposed for the new trajectory at this time were that all six phase-space variables should equal those of the Fig. 1 trajectory at the same time t'_I, except that we set $P'(t'_I) = P(t'_I) - 300$. This displacement of P, while keeping all else fixed, produced bound motion, as in the Fig. 1 trajectory, but with higher amplitude oscillation in the local well of $V(q,t)$; hence higher S, and so earlier stalling time $t_S \doteq 43.7$.

The dashed trajectory of Fig. 4 has only slightly higher E_5 than the solid trajectory; but this small difference eventually makes a large difference. During the driving phase of motion, the dashed trajectory of $Q(t)$ as shown in Fig. A3a) does not visibly differ from the original one, because the difference is only in oscillations to small to be seen on the plotted scale. The slightly higher E_5 in the dashed case makes a substantial difference in the amount of work the daemon does, however, because the higher E_5 of the dashed trajectory means higher S and earlier t_S, so that steady downconversion stops sooner in the dashed trajectory, as we see in both parts of Fig. A3, as well as in the main text's Fig. 4. In the dashed case the daemon does nearly 39% less total work on the slow mass than in the original trajectory of Fig. 1. The general limitation on daemon efficiency from the adiabatic invariance of S is not crippling, but neither is it trivial.